Journal of Environmental Economics and Management 89 (2018) 255–277

Contents lists available at ScienceDirect

Journal of Environmental Economics and Management

journal homepage: www.elsevier.com/locate/jeem

Anticipation and environmental regulation[☆]

Katherine Rittenhouse[a], Matthew Zaragoza-Watkins[b,*]

[a] Department of Economics, University of California, San Diego, United States

[b] Department of Economics, Vanderbilt University, 322 Calhoun Hall, Nashville, TN 37240, United States

ARTICLE INFO

Article history:
Received 3 October 2017
Revised 19 January 2018
Accepted 18 March 2018
Available online 6 April 2018

JEL:
C41
D92
Q58

ABSTRACT

When agents expect a change in regulation to change the relative price of new durable goods they may shift purchases forward to avoid compliance costs. In the context of new-vehicle emission standards, prior analyses have not considered this adjustment margin. We model the effects of anticipation on sales and retirements of durable goods, and test our theory's predictions empirically using the 2007 implementation of heavy-duty emission standards. We find evidence that anticipation caused a sales spike just before the policy took effect and a symmetric sales slump after implementation, which resulted in 31,164 more freight-truck sales ahead of the new standard and as much as $118 million in environmental damages over the lifetimes of those vehicles.

1. Introduction

In the United States, the transportation sector is the largest source of criteria pollution, and the second largest source of greenhouse gas (GHG) pollution, regulated under the Clean Air Act (CAA).[1] To help achieve NAAQS, EPA regularly updates pollution control standards for mobile and stationary sources. In recent years, concerns about local air pollution, global climate change, and energy security have prompted EPA to adopt more stringent standards for new mobile sources, requiring Original Equipment Manufacturers (OEMs) to deploy control technology, at some cost.

Vintage-differentiated emission standards are important and frequently-used tools for controlling pollution (Stavins, 2006). However, they are not theoretically optimal policy instruments[2]; generally, these regulations do not target pollution from all

☆ We are grateful for helpful comments from Hunt Allcott, Severin Borenstein, Lucas Davis, Gloria Helfand, Larry Karp, Michael Mehling, Reed Walker, Catherine Wolfram and two anonymous referees. We thank seminar participants at EDF, Harvard, MIT, Vanderbilt, UC Berkeley, and U.S. EPA, and participants at the MIT CEEPR workshop and the AERE Summer Conference for helpful comments. We also thank MIT CEEPR for financial support in acquiring data. All remaining errors are our own.

* Corresponding author.

E-mail addresses: krittenh@ucsd.edu (K. Rittenhouse), matthew.zaragoza-watkins@vanderbilt.edu (M. Zaragoza-Watkins).

[1] The CAA requires the U.S. Environmental Protection Agency (EPA) to set National Ambient Air Quality Standards (NAAQS) for six categories of common air pollutants. These "criteria" pollutants, which are so named because EPA is required to set NAAQS based on scientifically-supported health criteria, are particulate matter (PM), ground-level ozone (O_3), carbon monoxide (CO), sulfur oxides (SO_x), nitrogen oxides (NO_x), and lead (Pb). Of criteria pollutants, PM and ground-level O_3 present the broadest and most acute threats to human health and the environment. Transportation is also the largest source of volatile organic compounds, which, along with NO_x, are precursors to photochemical O_3.

https://doi.org/10.1016/j.jeem.2018.03.005
0095-0696/

sources or through all emission pathways.[3] When cost increases are unaccompanied by offsetting increases in private benefits, new-vehicle emission standards prompt motorists to operate their older, and higher-emitting, vehicles for longer than they otherwise would have, slowing fleet turnover and degrading the short-term environmental benefits of the policy (Gruenspecht, 1982).[4] Discrete changes in regulation also directly influence the timing of new-vehicle sales. Forward-looking consumers who wish to avoid paying the incremental cost of complying with a new standard may choose to "pre-buy" a new vehicle during the period just before a regulation is implemented. While the incentive for consumers to purchase new vehicles ahead of regulation may be intuitively clear, the impacts of discrete policy changes on new-vehicle sales and the composition of the fleet, and the implications for the environmental effectiveness of these policies are essentially unstudied.[5]

To investigate the incentives created by time-based notches in new-vehicle standards, the impact of those incentives on the new-vehicle sales cycle, and the implications for the effectiveness of new-vehicle standards, we analyze the market for new Class-8 heavy-duty vehicles (HDVs or trucks). We choose trucks for several reasons. They are important to the economy; the trucking industry carried 71% of the value and 70% of the weight of U.S. freight in 2007 (U.S. Department of Transportation, 2010). Heavy-duty trucks are responsible for a significant portion of the nation's air pollution from transportation. EPA has developed and implemented multiple rounds of new-truck emission standards, creating repeated opportunities to analyze the short-run impact of standards on new-truck sales. The trucking industry also has useful properties for analyzing strategic purchasing behavior. New trucks are relatively homogeneous durable goods, which are elastically supplied by competitive OEMs. Decisions about when to buy and retire trucks are generally made by firms seeking to maximize profit. There is relatively free entry into the freight transport market, which is where trucks are primarily used.

We address four specific questions: How does anticipation of a discrete change in regulation affect the pattern of new-truck sales? How does this change in the pattern of new-truck sales affect the pattern of used-truck retirements? How do these changes in purchasing and retirement patterns affect the environmental benefits of standards? Empirically, have recent discrete changes in regulation caused firms to pre-buy trucks?

We begin by developing a dynamic model of a competitive durable goods market, where firms incorporate upfront costs, operating costs and operating revenue into their purchasing and retirement decisions. We describe our model in the context of the HDV freight market. Operating costs of individual trucks are assumed to monotonically increase with age, motivating cycles of retirement and replacement. We derive necessary conditions for trucks to enter and retire from the freight market, and calculate comparative statics for changes in upfront and operating costs. As in many models of capital turnover in competitive markets, trucks enter when the net present value (NPV) of future operating profits equals or exceeds the new-truck price, while trucks exit as soon as operating costs exceed revenues. We find that an increase in the upfront cost of new trucks causes an increase in the equilibrium freight rate and vehicle lifetime (consistent with Gruenspecht), while an increase in the operating cost of new vehicles causes an increase in the equilibrium freight rate, but has an ambiguous effect on vehicle lifetime.

We then analyze how incorporating anticipation (i.e., beliefs about future new-truck prices) affects investment and retirement patterns. In particular, we consider how an anticipated increase in the new-truck price affects pre-regulation purchasing decisions (holding constant the new-truck price in the pre-regulation period).[6] In some contrast to models without anticipation, where higher upfront costs initially deter entry (allowing used durables to earn rents), forward-looking purchasers (chasing the opportunity to earn future rents) buy trucks just ahead of the regulation (fully dissipating those rents). That is, in the period

[2] Vintage-differentiated emission standards may be cost-effective in the short run, as retrofitting existing units is often more expensive than requiring that new units be built to a certain standard. However, these regulations create an incentive for operators of older, higher-emitting units to delay retirement. This dynamic inefficiency erodes long-term cost-effectiveness, and may reduce the environmental effectiveness of the standard. In some cases, standards may even backfire if increased lifetime emissions from older units outweigh emissions reductions from new units. That these policies cause a delay in fleet turnover is seen theoretically and empirically in multiple fields of environmental regulation. Two prominent examples include regulation of power generators under the Clean Air Act's New Source Review program (Maloney and Brady, 1988), and automobile emissions, safety, and fuel-economy standards (Gruenspecht, 1982; Jacobsen and Van Benthem, 2015).

[3] Optimal pollution control could be achieved by imposing a series of pollution taxes equal to the marginal social damage from emitting a quantity of the relevant pollutant. Pigouvian taxes would induce vehicle operators to drive fewer miles and buy less-polluting vehicles, OEMs to market vehicles with better pollution control equipment, and refiners to produce cleaner fuels (Fullerton and West, 2002). Capturing each of these adjustment margins would yield the broadest and most cost-effective suite of emissions reductions.

[4] Known as the "Gruenspecht" effect, this phenomenon may be particularly relevant for criteria pollutant standards, where the regulation increases production costs for OEMs, and thus private costs to vehicle purchasers, while socializing the benefits (i.e., cleaner air). This is in contrast to, for example, GHG and fuel-economy standards, which provide significant direct consumer benefits in the form of fuel savings, resulting in lower operating costs.

[5] A discrete change in regulation is an example of a policy "notch". The drawbacks of policy notches are well studied (see, for example, Sallee and Slemrod (2012)). Our work is most closely related to a recent literature on time-based tax-policy notches, which examines the effect of anticipated changes in sales and consumption taxes on purchase patterns (see, for example, Crossley et al. (2014), Cashin and Unayama (2016), and Coglianese et al. (2016)). A related, but distinct, literature investigates the impacts of short-lived fiscal stimulus (see, for example, D'Acunto et al. (2016)). In particular, Mian and Sufi (2012), Li et al. (2013), and Hoekstra et al. (2017) study the impact of the 2009 *Cash for Clunkers* program. These studies each find that the program initially induced the purchase of several hundred thousand vehicles, but that fewer vehicles were sold in the months after the program ended, completely eroding the short-run sales gains. These empirical results are consistent with the predictions of our theory model, but the policy change and setting we study, and the market dynamics we explore, are distinct. With the exception of Lam and Bausell (2007), which we discuss below, ours is the first analysis of the influence of "anticipation" on the effectiveness of emission standards.

[6] In our model, trucks are elastically supplied by a competitive industry. Thus, the expected price increase does not create an incentive for OEMs to increase prices in the current period. In practice, firms may be able to exert some market power. We abstract from imperfect competition in order to focus the analysis on demand side effects. Modeling imperfect competition would not qualitatively change our results. If firms were choosing quantities in response to a downward sloping demand curve, the surplus accruing to vehicles purchased ahead of regulation would still increase, and the share of that increase that flows to vehicle purchasers would induce a pre-buy.

immediately preceding regulation, firms add trucks to the fleet until the anticipated windfall to pre-regulation capital is fully arbitraged.[7]

Strategic pre-buying affects retirements and post-regulation new-truck sales. Injecting new trucks into the freight transport market lowers the current period freight rate, reducing the value of used trucks and causing the oldest trucks to immediately retire. The lower freight rate also extends the post-regulation period during which firms are discouraged from investing in new trucks, because long-run profits are too low to compensate for the higher purchase price. Though pre-buying distorts the transition to the new equilibrium, the predicted effect of anticipation on long-run sales is zero. That is, the increase in pre-regulation purchases induced by rent seeking is predicted to equal the decrease in post-regulation sales, causing no change in the long-run quantity of vehicles sold.[8]

The net environmental effect of anticipation depends on how gains from accelerated turnover compare with losses from more-modest emission-rate improvements. If older trucks are higher emitting, accelerating retirement provides an immediate environmental benefit. Post-regulation, deterred investment reduces the environmental effectiveness of the policy.

To empirically test our predictions, we estimate a model of new-truck sales, using monthly Class-8 truck sales in the U.S. over the period 1991–2015. Our preferred specification includes the monthly average real oil price and quarterly U.S. Gross Domestic Product (GDP), as well as year and month-of-year fixed effects. This model explains the vast majority of variation in truck sales. Oil prices and GDP are significant drivers of truck sales, both of which were changing rapidly and dramatically around the time these standards were implemented. Our econometric approach attempts to disentangle the impact of anticipation from those of other sales drivers. We investigate whether anticipation affected sales by examining residual variation in sales (i.e., variation not explained by our model) around the month the standards took effect. If anticipation did impact the sales cycle, our model would predict persistently positive residuals just before the policy start date and persistently negative residuals immediately after the policy start date.

During this time, EPA implemented four rounds of new-truck criteria pollutant standards. We focus on the 2007 new-engine standards, which are widely regarded as the most significant regulatory action (i.e., with respect to trucks) taken by EPA during the 25-year span of our data. Consistent with our prediction, we find robust evidence of a sales spike in the months before implementation of the 2007 criteria pollutant standards, followed by a similarly-sized sales dip in the months after the regulation took effect. Across various specifications, we estimate anticipation of the 2007 criteria pollutant standards caused several thousand more trucks to be sold in each of the months prior to, and approximately the same number fewer trucks to be sold in each of the months after, the introduction of the standards, resulting in a net sales impact (of anticipation) which is statistically indistinguishable from zero.

Our study makes several contributions to the literature on standard-based (environmental) regulation. Our theoretical model introduces the concept of anticipation, distinguishing between the effects of this form of strategic behavior, the direct effect of policy on new-vehicle sales, and the indirect (i.e. Gruenspecht) effect on fleet composition. Our model illuminates two offsetting channels through which anticipation impacts the environmental effectiveness of standards, the net effect of which is *a priori* ambiguous. Empirically, we confirm that strategic responses to regulation affect investment cycles. Our results have important implications for policy design and program evaluation. Confounding the effects of anticipation with the direct effects of policy would, under a variety of identification strategies, result in significantly biased estimates.[9] As noted in the growing literature on tax avoidance, it is critical that analysts account for this behavior when studying markets in which agents can shift the timing of purchases in anticipation of new regulation.

The paper proceeds as follows. In Section 2 we provide a brief background on criteria pollutant regulations for HDVs in the U.S. Section 3 presents our theoretical model of the truck market, and explores the impacts of anticipated and unanticipated regulation on the cycle of capital turnover and the associated emissions. Section 4 describes the econometric framework and Section 5 the data used to estimate the effect of anticipation on the pattern of new-truck sales. In Sections 6 and 7 we present and discuss our results and Section 8 concludes.

2. Background

EPA has regulated emissions from HDVs for over forty years (U.S. EPA, "Heavy-Duty Highway Compression-Ignition Engines and Urban Buses – Exhaust Emission Standards"). In 2001, EPA finalized criteria pollutant emission standards for diesel heavy-duty engines and vehicles, model-year 2007 and later, significantly tightening the standards for PM and NO_x emissions

[7] The pre-buy is predicted to be strongest during the period just before regulation, because earlier purchases earn fewer revenues over the life of the truck. Strategic timing of purchases is explored in more depth in our theory section.

[8] This effect of anticipation is distinct from the direct effect of the policy on sales, which, as noted above, is predicted to be a decline in sales following the introduction of regulation.

[9] For example, in the absence of strategic behavior, one might use a regression discontinuity design to determine the impact of standards on sales, measuring the difference in sales before and after the regulation was implemented. However, given the predictions of our theory model, a regression discontinuity approach would likely produce biased results. In particular, if forward-looking agents strategically move the time of their purchase ahead of regulation, comparing vehicle sales prior to implementation with sales after implementation would bias results towards showing a larger reduction in net sales and economic cost, in terms of lost production and employment.

(40 C.F.R. § 69, 80, 86).[10] For engines, the standard for PM decreased from 0.1 to 0.01 g per brake-horsepower hour (g/bhp-hr), NO_x standards decreased from 2.4 to 0.2 g/bhp-hr, and standards for non-methane hydrocarbons (NMHC) were put in place (in isolation from NO_x standards, for the first time) at 0.14 g/bhp-hr. Standards were similarly tightened for overall vehicle performance, varying with gross vehicle weight class. EPA projected that the new standards would reduce annual NO_x, NMHC, and PM emissions by 2.6 million, 115,000 and 109,000 tons respectively, preventing an estimated 8300 deaths, 9500 hospitalizations, and 1.5 million lost workdays (40 C.F.R. § 69, 80, 86).

The 2007 criteria pollutant standards were expected to increase the total cost of purchasing and operating model-year 2007 and later HDVs. Though the price for larger HDVs can extend well beyond \$100,000 per vehicle (40 C.F.R. § 69, 80, 86), it was generally accepted that the cost of the 2007 criteria pollutant standards would represent a significant increase in the cost of new HDVs. EPA projected upfront costs for HDVs would increase by \$3230 in early years, and lifetime operating costs would increase by \$4600 (EPA, 2000). Industry analysts and OEMs predicted that surcharges of between \$7000 and \$10,000 would *need* to be applied to model-year 2007 and later HDVs, to cover the cost of installing pollution-control technology. The discrepancy between EPA and industry projections of compliance costs likely contributed to uncertainty among HDV fleet operators. Leading up to the 2007 criteria pollutant standards, the possibility of a pre-buy was widely discussed in the popular press, while some industry analysts advised purchasers to pre-buy HDVs.[11]

3. Theory

3.1. A model of capital turnover

To analyze how anticipation of future price changes impacts the sales patterns and fleet composition in the HDV market, we develop a theory model, where profit-maximizing agents choose when to purchase and retire vehicles which they operate in a competitive freight market. Note, although we describe our model in the context of the HDV market, it applies more generally as a capital turnover model. The predictions of our theory model hold whenever agents anticipate discrete and costly changes to durable goods. We begin by describing the lifetime profit associated with a vehicle in this market:

$$\Pi = \int_0^T (P(Q(t)) - C(t))e^{-rt}dt - M. \tag{1}$$

Where: T is the lifetime of the vehicle; $P(Q(t))$ is the freight rate (i.e., operating revenue) at time t, which, fixing demand, is a function of the aggregate supply of vehicles in the market; $C(t)$ is the operating and maintenance cost for a vehicle at time t, which increases monotonically as t approaches T; r is the discount rate; and M is the purchase price of a new vehicle.

A price-taking firm's profit maximization problem can be thought of in terms of its vehicle purchase and retirement decisions. A firm will choose to purchase and operate a new vehicle if the NPV of operating that vehicle is greater than the purchase price:

$$\int_0^T (P(Q(t)) - C(t))e^{-rt}dt \geq M. \tag{2}$$

A firm will continue to purchase and operate new vehicles until the point where the profit gained from adding the last vehicle equals zero:

$$\int_0^T (P(Q(t)) - C(t))e^{-rt}dt = M. \tag{3}$$

A firm chooses when to retire a vehicle based on the first-order condition for profit maximization with respect to T:

$$\max_T \int_0^T (P(Q(t)) - C(t))e^{-rt}dt - M. \tag{4}$$

Differentiating with respect to T, we get the first-order condition for profit maximization:

$$P(Q(T)) = C(T). \tag{5}$$

Thus, a price-taking firm retires a vehicle at the time, T, when revenue equals costs.

When a truck is retired it is immediately replaced by a new vehicle. It follows directly from our specification of $C(t)$ that the freight-services supply curve is the aggregate cost function of all vehicles in service. In equilibrium, both supply of freight

[10] The regulation states the new standard will bind on model-year 2007 vehicles. However, unlike light- and medium-duty vehicles, freight trucks do not follow model-year production cycles. They are produced in response to orders received from fleets and individual purchasers. Leading up to the 2007 change in HDV standards, lag times between orders and deliveries increased (in response to the spike in demand). When it came time to apply the new standard, this created a potentially gray area for enforcement, as the regulation does not specify whether the new standards should be enforced based on the order date or sale (i.e., delivery) date. EPA technical staff confirm that, in order to allow OEMs time to complete orders made in 2006, the new standards were enforced on vehicles with sale dates after June 30, 2007. For ease of exposition, we refer to vehicles delivered after the change in enforcement as model-year 2007 vehicles.

[11] See, for examples, "Emission Rule Change the Engine for Truck Sales: But this year's boom to be next year's bust." Chicago Tribune. Web. 23 May 2016., and "To pre-buy or not." Fleet Owner. Web. 23 May 2016., respectively.

services and demand for those services are constant over time, yielding a constant quantity, Q, of vehicles in the market, and price, P, for the services those vehicles provide.

3.2. The equilibrium effects of regulation

We are interested in the effect of new-vehicle emission standards on the HDV market. Emission standards may affect purchase and retirement decisions through two channels: they may cause OEMs to install additional abatement technology, increasing the purchase price of a new vehicle; and they may change the operating-cost function of a new vehicle. Without loss of generality, we model the effect of a change to the purchase price.[12]

Suppose the implementation of a new-vehicle standard increases the purchase price, M, but does not change the cost function. How will this change affect P, Q and T? Note that a change in M will directly affect a firm's entry decision (Equation (3)) without directly affecting the exit condition for existing vehicles (Equation (5)). We begin by evaluating the comparative statics on the entry condition, and find that the equilibrium freight rate is increasing in M. That is, P is increasing in M. Fixing demand, it follows that a greater equilibrium P implies a lower equilibrium Q, while the exit condition implies a longer vehicle lifetime, T. These equilibrium comparative statics are analyzed in Appendix A.

Remark. Our model confirms two distinct and previously identified effects of an unanticipated price change. First, a price change affects the flow of new vehicles into the market: the quantity of vehicles purchased during a given timeframe is inversely related to the purchase price of the vehicle. We refer to this as the direct effect of the price change. Second, a change in the purchase price affects the stock of vehicles operating in the market: the quantity of vehicles retired during a given timeframe is inversely related to the purchase price of a new vehicle. Equivalently, lifetimes for all vehicles in the fleet are increasing in the new-vehicle purchase price. We refer to this as the Gruenspecht effect.[13] Foreshadowing the discussion below, the existence of the Gruenspecht effect suggests that an increase in the flow of vehicles entering the market which reduces the NPV of used vehicles will drive an increase in retirements.

3.3. Transitions and the effects of anticipation

We have established that regulation which changes the purchase price of new vehicles results in a new equilibrium freight price, quantity, and vehicle lifetime. We now turn to the question of how the market transitions between equilibria. We first assume firms do not anticipate regulation. This is the standard form of analysis in the literature. We then examine how anticipation, through its impact on the transition, affects the pattern of new-vehicle sales and the fleet composition. Without loss of generality, we analyze the case of an increase in M, from M_1 to $M_2 > M_1$, holding $C(t)$ constant.

3.4. Unanticipated regulation

We first consider the case where the price increase is unanticipated, that is, firms respond to a change in price only after it has occurred. An increase in M does not immediately impact either $C(t)$ or P. The exit condition is met, and the first retirement after the price change occurs as it otherwise would have. However, that retired vehicle will not be replaced by a new vehicle purchase, as it would have been in equilibrium. The new entry condition will not be met until the NPV of a new vehicle increases to equal the new purchase price, M_2. As vehicles retire without replacement, the supply curve shifts inward (or upward). As the fleet ages, the quantity of vehicles in service declines and the freight rate rises, increasing the NPV of vehicles in the market. When enough vehicles have retired such that the new entry condition is met, the market reaches its new equilibrium quantity and price, Q_2 and P_2, respectively. Once this new equilibrium is reached, each retired vehicle is again replaced by a new purchase. The exit condition becomes $C(T_2) = P_2$, giving vehicles a longer lifetime T_2 for a higher P_2. Empirically, we would expect the period immediately following the price change and during the transition to a new equilibrium to appear as a sales slump.

During this transition period, and once the new equilibrium is reached, the freight rate is greater than P_1. Accordingly, vehicles purchased prior to the price change are earning greater revenue than expected at the time of their purchase. A vehicle purchased immediately before the implementation of new standards therefore accumulates a windfall for its entire lifetime, earning:

$$\int_0^{\tau} (P(Q(t)) - C(t))e^{-rt}dt + \int_{\tau}^{T_2} (P_2 - C(t))e^{-rt}dt. \tag{6}$$

[12] In Appendix A, we also consider the effect of a change to the cost function, showing that for each change in M there is an analogous change in $C(t)$ which will have exactly similar effects on equilibrium P and Q.

[13] Gruenspecht (1982) presents a structural model of vehicle fleet turnover under vintage-differentiated emission standards. As emission standards become more stringent, lifetime (upfront and possibly maintenance) costs of new vehicles subject to those standards rise. Gruenspecht shows that an increase in the price of new cars induces substitution towards used cars, which increases the value of used vehicles. In Gruenspecht's model, the scrappage rate depends on the value of used vehicles, where the scrappage condition is that the value is below a threshold. Thus, increasing the value of used vehicles drives down scrappage rates for existing vehicles, reducing the environmental effectiveness of the standard. Our model assumes a zero scrappage value, making Gruenspecht's scrappage rate analogous to our retirement rate.

Here, price, $P(Q(t))$, is increasing from P_1 to P_2 between $t=0$ and $t=\tau$, after which it is steady at P_2 for the remainder of the vehicle's lifetime, until retirement at T_2. Rents associated with the difference in expected and earned profits due to the unanticipated price change are characterized by:

$$Rents = \int_0^{\tau} (P(Q(t)) - P_1)e^{-rt}dt + \int_{\tau}^{T_1} (P_2 - P_1)e^{-rt}dt + \int_{T_1}^{T_2} (P_2 - C(t))e^{-rt}dt. \tag{7}$$

Looking at the right-hand side of this equation, the first integral represents the difference in the earned and expected revenues during the period when the market is in transition to a new equilibrium. The second integral represents the difference in earned and expected revenues from the time the market reaches a new equilibrium until the time at which the vehicle would have been retired, T_1. The third and final integral represents the revenue earned during the period of extended lifetime, from the time T_1 which the vehicle was expected to retire when purchased, until the time T_2 when the vehicle is retired under the new equilibrium price. Each of these integrals is positive, and increasing in P_2. In a world where firms anticipate, and can respond to future changes in M, the prospect of earning these rents provides an incentive to purchase additional vehicles before the price changes.

3.5. Anticipated regulation

In reality, new regulations are implemented through an often-lengthy public process, which provides regulated entities with knowledge about the timing and content of forthcoming standards. When firms anticipate a change in the price of new vehicles, and foresee the rents associated with a higher future equilibrium freight rate, we expect them to arbitrage those rents by adding vehicles to the market. This pre-buy immediately lowers the freight rate, driving down revenues accruing to vehicles in the market. Firms will continue to add vehicles until the point where the negative rents earned when $P(Q(t)) < P_1$ equal, in absolute value, the positive rents earned when $P(Q(t)) > P_1$:

$$\int_0^{\tau_0} (P(Q(t)) - P_1)e^{-rt}dt = \int_{\tau_0}^{\tau} (P(Q(t)) - P_1)e^{-rt}dt + \int_{\tau}^{T_1} (P_2 - P_1)e^{-rt}dt + \int_{T_1}^{T_2} (P_2 - C(t))e^{-rt}dt. \tag{8}$$

The left-hand side of this equation represents the period from the vehicle purchase until the time $t=\tau_0$, when the market price $P(Q(t))$ reaches the price at the old equilibrium, P_1. During this time, vehicles earn lower revenues than they would have under original equilibrium, giving the integral a negative value. This offsets the higher rents earned as $P(Q(t))$ rises above P_1 for the remainder of the vehicle's lifetime. The first right-hand side integral represents the time between $t=\tau_0$ and $t=\tau$ when the market price is above the original equilibrium price, P_1, but below the new equilibrium price, P_2. The second and third right-hand side integrals correspond to the second and third integrals in Equation (7), above. These three right-hand integrals are positive, and their sum is equal, in absolute value, to the left-hand side integral.

This equality implies a larger pre-buy for a larger price change: The higher is P_2 relative to P_1, the larger are the rents on the right-hand side of this equation, the more negative must be the rents on the left-hand side and, thus, the larger in volume must be the pre-buy.[14]

Empirically, we expect to see an increase in vehicle sales directly prior to regulation, followed by a sales slump directly after regulation is implemented. Note, this sales slump is distinct from and additional to the direct effect of the policy on new-vehicle sales.

Anticipation also affects the composition of the fleet through the exit condition. Recall that a firm will retire a vehicle at the time, T, when its cost, $C(T)$, equals the freight rate, $P(Q(T))$. When firms anticipate a price change, the injection of pre-bought vehicles immediately drives down the freight rate below P_1. Those vehicles whose costs are greater than the new freight rate are no longer profitable to keep in operation, and will be pushed out of the market, into early retirement.[15] Importantly, the quantity of vehicles pushed out must be less than the quantity of pre-bought vehicles. Only with an increase in total quantity will price decline in the short term, dissipating all available rents. We will return to this observation when considering the impact of the pre-buy on emissions.

Remark. Note that the new steady-state quantity and price are the same whether or not the price change is anticipated. This implies that the effect of anticipation on vehicle flows must be a symmetric shift in vehicle sales.

[14] We treat the time between when the price change is announced and when it takes effect as a single decision point. In practice, firms often know many months in advance that a price change will occur. One question this raises is whether one would expect to see a pre-buy well in advance of the price change, as firms seek to capture the future stream of rents. While we do not explicitly model this tradeoff, we do consider the following thought experiment. Suppose a firm chooses to purchase a vehicle just ahead of when all other "pre-buy" vehicles are purchased. This vehicle would earn the same stream of revenues as the other "pre-buy" vehicles, except that it would initially earn some $P' < P_1$ (i.e., because it entered, pushing down the price) and not earn P_2 for as long an interval at the end of life. If the average flow of revenues to the other "pre-buy" vehicles is P_1 – which must be the case in order to satisfy the entry conditions with equality – then subtracting a portion $P_2 > P_1$ and replacing it with $P' < P_1$ must yield lower revenues, and therefore be dominated. In contrast, imperfect competition among manufacturers could generate an equilibrium, where sellers increase prices each period between when the price change is announced and when it takes effect. In this setting, purchasers would execute a mixed strategy, trading off reduced rents associated with making an earlier purchase against obtaining a lower upfront price.

[15] Under changing prices, this exit condition becomes more complex, as there may be multiple points where $C(t) = P(Q(t))$, multiple local maxima. A vehicle will exit the market at the time T which produces the global maximum profit.

Remark. Our model predicts that anticipation of a price change affects both the flow and stock of vehicles. These effects are additional to the previously identified direct and Gruenspecht effects, respectively. Anticipation of costly regulation increases the flow of vehicles into the market directly prior to regulation implementation, and symmetrically decreases the flow of compliant vehicles post-implementation, reducing the environmental effectiveness of the policy. Through the same channel as the Gruenspecht effect, anticipation which increases the flow of vehicles ahead of regulation implementation is expected to reduce the remaining lifetime revenues and thus the value of used vehicles. This pushes the oldest, highest-emitting vehicles in the fleet to an early retirement, increasing the environmental effectiveness of the policy. Note, the net effect of anticipation on the emissions impact of the policy is *a priori* ambiguous. The net emissions impact will depend on the differences in the emission rates of the three categories of vehicles, and the shape of the supply and demand curves. We relegate our discussion of the emissions impact of anticipation to Appendix A.

4. Econometric framework

In this section we present our approach for measuring the effect of anticipation on the new-vehicle sales cycle. Our setting is the 2007 implementation of criteria pollutant standards for HDVs, which is widely regarded as the most significant regulatory action taken by EPA with respect to trucks during the span of our data.[16] Given the expected increases in purchase price and operating costs, our model predicts a sales spike directly prior to implementation, followed by a sales slump in the months after the regulation took effect.

In choosing an econometric framework to test the predictions of our theory model, we face several empirical challenges. First, the timing of anticipation's effect on the sales cycle is not immediately obvious. While our model predicts firms will pre-buy in the period immediately before the standards are implemented, concerns around vehicle availability may cause some to shift their purchases further ahead of regulation. In addition, GDP and oil prices – two key drivers of HDV demand – were moving around the time the policy was implemented. Failing to account for these drivers would likely bias our estimates.[17] Finally, we do not find a well-suited comparison population to serve as our counterfactual, so we must rely on time-series variation to identify the effect of anticipation.[18]

To address these empirical challenges, we exploit our twenty-five-year time-series to estimate a model of monthly HDV sales. Similar to Lam and Bausell (2007), we model monthly U.S. HDV sales, Q_t, as a linear function of several economic drivers.[19]

$$Q_t = \beta_0 + \beta_1 GDP_t + \beta_2 Real_t + \beta_3 X_t + \varepsilon_t \tag{9}$$

Here, GDP_t and $Real_t$ are GDP in the quarter corresponding to month t and the real price of oil in month t, respectively. X_t is a set of year and month-of-year binary variables equal to one in their designated year or month-of-year, respectively, and zero, otherwise. ε_t is the econometric error term, which accounts for period-specific sales shocks that are orthogonal to other specified sales drivers.

GDP is an important indicator of the strength of the economy, which is used in many industries, including trucking, to forecast the strength of future demand for their products and services. We expect truck operators to purchase more trucks in response to strong economic growth, which implies a positive β_1 in Equation (9).

We use real oil price as a proxy for operating costs.[20] Fuel is the largest single cost for trucking fleets, accounting for 38% of the cost of ownership. Lease or purchase payments, by contrast, account for 10% of total average operating costs to owners (Torrey and Murray, 2014). Moreover, other operating costs (e.g., maintenance) are likely correlated with implementation of the regulation. Including them could bias our estimates (Lam and Bausell, 2007). According to our theory model, higher

[16] EPA and industry analysts differed in their predictions of the associated cost increase. EPA proposed the new equipment would add $3200 to the cost of each new truck, while industry claimed the cost would be closer to $10,000. The uncertainty associated with cost increases and reliability of new technologies led some industry analysts to recommend pre-buying in response to the new standards.

[17] This *approach* is taken in several industry-funded analyses of the 2007 implementation (see, for examples, Harrison and LeBel (2008) and Calpin and Plaza-Jennings (2012)).

[18] We consider and reject three comparison groups: new-HDV sales in Canada, new-HDV sales in Mexico, and new-light duty vehicle (LDV) sales in the U.S. Canada adopted the same HDV standards simultaneously with the U.S., and Mexico adopted HDV standards for the following model-year, within the plausible treatment window of the U.S. policy. Further, new-vehicle sales in Canada and Mexico may have been impacted by U.S. HDV policy, violating the stable unit treatment value assumption. Finally, while LDV sales likely share some drivers with HDV sales (including oil price and the state of the macroeconomy), it is unlikely that those drivers affect HDV and LDV sales in the same way, making the comparison unsuitable for panel-data identification strategies.

[19] Lam and Bausell (2007) attempt to measure the impact of EPA's 2002 HDV-engine standards on the timing of truck production, providing the only prior empirical evidence of pre-buying in response to vehicle emissions standards. They specify an econometric model of truck production as a function of GDP, diesel fuel prices, average retail truck prices, quarter-of-year fixed effects and a linear time trend. Lam and Bausell do not consider the effect of anticipation on sales after the implementation of the regulation. To estimate the size of the pre-buy, they use a binary variable that takes a value of one in the six months prior to the implementation of the 2002 regulation, and zero otherwise, and find that production increased by approximately 20% in the six months prior to implementation. Because anticipation has an echo in the post period, we suspect their estimation strategy produces biased results.

[20] While diesel prices may be a more direct measure of fuel costs, we find that oil prices are a stronger predictor of vehicle sales, particularly in the presence of month-of-year fixed effects. This appears to be due to a spurious correlation between the annual cyclicality of vehicle sales and diesel fuel prices (likely driven by seasonal variation in home-heating-oil demand). Additionally, the global oil price is less likely to be affected by local shocks, which may affect both diesel price and HDV sales and which would potentially bias our estimates of the effect of diesel fuel costs on sales.

operating costs will reduce lifetime profit, reducing demand for new freight trucks, so we expect β_2 in Equation (9) to be negative.[21]

Month-of-year binary variables are included to flexibly model regular variation in truck sales. These monthly differences are likely due to the cyclicality of new model-year releases. Yearly binary variables are included to flexibly model average differences in sales across years, which may be due to differences in year-specific product attributes (and the level of regulation).

Vehicle purchase price is excluded from the estimating equation. The pollution-control technology required to comply with the 2007 standards increased the cost of producing a new HDV. If these costs were passed on to HDV purchasers, or if price was endogenous to demand, including price in the regression would absorb a portion of the sales variation caused by the regulation. Additionally, we don't find a reliable instrument for vehicle supply.[22] If anticipation drove manufacturers to change prices ahead of regulation (i.e., increase prices in response to a positive demand shock), then excluding price from the estimating equation would attenuate our estimates of the effects of anticipation on new-vehicle sales. To the extent that vehicle prices vary systematically by year or month-of-year (e.g., with the release of new model-year vehicles), the fixed effects should absorb the resulting sales variation.

The econometric error – variation in the data not explained by the econometric model – provides our measure of anticipation. Under the conditional unconfoundedness assumption, predictions from this model provide valid counterfactual estimates of monthly HDV sales. Under the added assumption that HDVs are elastically supplied, differences between predicted and observed monthly HDV sales (i.e., the residuals from our sales model) can be interpreted as demand shocks. If the model fits the data, then residuals will vary quasi-randomly around zero. To identify the effect of anticipation, we investigate the pattern of sales residuals for evidence of demand shocks around the time the regulation was implemented: persistently positive residuals just before the policy start and persistently negative residuals after the policy start.[23] An advantage of our approach is that it does not require any strong assumptions about the time period in which anticipation might have affected HDV sales. A drawback of our approach is that it does not allow us to identify the direct effect of the policy change on the long-run level of HDV sales, which would be necessary for a fuller welfare analysis, or to quantify the magnitude of the bias that derives from failing to account for anticipation.[24]

In addition to a graphical analysis, we test for the presence of anticipation using OLS. Our first approach is to regress the residuals from Equation (9) on two binary variables:

$$\varepsilon_t = Pre_t + Post_t + \epsilon_t. \tag{10}$$

Pre_t takes the value of one during the seven months prior to the regulation and zero otherwise, while $Post_t$ takes the value of one in the seven months after the regulation takes effect and zero otherwise. Both Pre_t and $Post_t$ take the value of zero for the month in which the regulation was implemented (July 2007).[25] ϵ_t is the remaining econometric error.

This approach implicitly supposes that Pre_t and $Post_t$ do not covary with the regressors from Equation (9). If they are (spuriously) correlated, excluding these binary variables from the main specification would produce biased coefficients. To address this, we also estimate a specification where Pre_t and $Post_t$ are included as regressors:

$$Q_t = \beta_0 + \beta_1 GDP_t + \beta_2 Real_t + \beta_3 X_t + Pre_t + Post_t + \varepsilon_t. \tag{11}$$

Our policy coefficients may yet be biased if we have incorrectly specified the relationships between sales and these drivers. One concern may be that the oil price was changing dramatically and unprecedentedly around the time of the policy. For robustness, we test the sensitivity of our results to alternative specifications of the relationship between oil prices and sales. Moreover, our policy coefficients will be biased if we have omitted sales drivers that covary with our policy variables. In particular, some readers may be concerned that the policy shortly preceded the "Great Recession". If macroeconomic sales drivers (e.g., expectations about the future demand for freight activity) were changing discretely around the time of the policy, our coefficients may be biased. To test the robustness of our results to these concerns, we estimate additional specifications, which include several

[21] Our investigation of the relationship between oil prices and the monthly new-truck producer price index (PPI) suggests new-truck prices are relatively insensitive to oil-price fluctuations. To the extent that oil price enters the new-truck supply function, we expect higher oil prices to drive higher new-truck prices, resulting in a reduced quantity demanded.

[22] Lam and Bausell (2007) instrument for the price of freight trucks with the price of steel, a significant input to truck production. However, closer examination of this relationship shows that the truck PPI is relatively insensitive to changes in the steel PPI. Further, the steel PPI is driven by changes in oil price, violating the exclusion restriction.

[23] The intellectual basis for our approach to accounting for anticipation comes, in part, from McCrary (2008). In the context of a regression discontinuity design – a method for identifying the effects of a treatment on individual outcomes when an arbitrary threshold along a continuous "running variable" determines treatment status – McCrary develops a test for manipulation of the running variable, whereby agents may affect their own treatment status. McCrary argues that endogenous manipulation of the running variable violates the assumption of continuity (i.e., quasi-random assignment around the treatment threshold) and thus invalidates a regression discontinuity identification strategy in these cases. Hausman and Rapson (2017) elaborate on the implications for the McCrary test when time determines treatment status. In these cases, the potential for endogenous movements in the outcome variable in response to the timing of treatment may present serious threats to identification. In our context, firms may easily manipulate the time of purchase in anticipation of costly HDV emissions standards, plausibly causing bunching of vehicle sales directly before the price change which would bias time-series estimates of the direct effect of the policy change on the long-run level of HDV sales.

[24] Including year fixed effects absorbs the direct effect of the policy change on the long-run level of HDV sales. Because the policy was effectively implemented mid year (i.e., July 2007), it should not influence our estimates of the effect of anticipation.

[25] The pattern of variation in Fig. 1, below, suggests that sales were affected between four and eight months on either side of the policy. We use seven months in our main specification, but test the sensitivity of our results to alternative treatment-period specifications in robustness checks, in Appendix B.

Table 1
Descriptive Statistics of Variables used in our Analysis.

Variable	Obs	Mean	Std. Dev.	Min	Max
HDV Sales	285	14367	4529	6232	26380
Oil Price (1991 $)	285	52.69	30.96	13.45	137.39
GDP ($ Billion)	285	11480	3429	6054	17600

Table 2
Demand regression results.

Variables	(1) Sales	(2) Sales	(3) Residuals	(4) Sales
Real Oil Price	−71.21	−54.21		−28.19
	(16.59)	(14.00)		(11.38)
GDP	0.646	4.330		5.018
	(0.153)	(1.610)		(1.295)
Pre-treatment			4030	4526
			(530.0)	(938.1)
Post-treatment			−4676	−5191
			(530.0)	(862.5)
Observations	278	278	278	278
R-squared	0.066	0.856	0.336	0.910
Year FE	N	Y	N	Y
Month FE	N	Y	N	Y

This table reports coefficients and standard errors for four specifications of our econometric model of truck sales. In Columns (1), (2) and (4), the dependent variable is monthly Class-8 HDV sales. Columns (1) and (2) present results from a specification which uses Real Oil Price and GDP as independent variables; the specification reported in Column (2) includes year and month-of-year fixed effects. In Column (3), the dependent variable is the difference in actual Class-8 HDV sales and Class-8 HDV sales predicted by the fixed-effects model presented in Column (2), and the independent variables are two binary variables which respectively take the value of one during the seven months prior to and the seven months after regulation takes effect. Column (4) presents results from a specification which includes these two binary variables in the fixed-effects specification presented in Column (2).

leading economic indicators and oil price futures contracts. Additionally, we test the sensitivity of our results to alternative specifications of treatment-period length. The details and results of these robustness checks are presented in Appendix B.

5. Data

Monthly Class-8 HDV sales for the U.S., from 1991 through 2015, were purchased from Ward's Automotive, Inc. Monthly real oil prices (in 1991 dollars) for the same time period were collected from the Energy Information Administration (EIA). Quarterly GDP data for the same period were obtained from Bureau of Economic Analysis (BEA), and matched to monthly observations in the same quarter and year. Table 1 presents descriptive statistics for our data.

6. Results

In this section we report results for our econometric model of sales, investigate residual variation not explained by our econometric model for evidence of anticipation using graphical and regression analysis, and test our identifying assumption.

6.1. Sales model

Table 2 reports coefficients and standard errors for several specifications of our econometric model of truck sales. Column (1) reports results for our base specification, excluding year and month-of-year fixed effects. While the coefficients on real oil price and GDP each take the predicted sign, and are highly statistically significant, they account for a relatively modest share of sales variation. To explain more of the variation, and to control for seasonality and year-specific disturbances that could confound our results, we add annual and month-of-year fixed effects. Column (2) reports results for this specification. Fixed effects significantly improve the fit of our econometric model. Coefficients on real oil price and GDP have the same signs as in Column (1) and are statistically significant at the one percent level. Adding fixed effects attenuates the impact of the real oil price on sales and amplifies the effect of GDP. We infer that the fixed effects are flexibly controlling for some variation in sales that covaries with, but is not driven by, the real oil price. At the mean, a one percent increase in GDP is associated with a four percent increase in HDV sales, while a one-hundred percent increase in the real oil price is associated with a twenty percent decline in sales.

NOTE: In this figure the y-axis reports the difference between actual monthly Class-8 HDV sales and monthly Class-8 HDV sales predicted by our fixed-effects regression, the coefficients and standard errors for which are presented in Column (2) of Table (2). The x-axis reports the date. Each dot is a monthly observation and the reference line is aligned with the month the regulation took effect (July 2007).

Fig. 1. Plot of monthly HDV sales residuals.

6.2. Graphical analysis of anticipation

Fig. 1 plots monthly residuals from our fixed-effects specification, the coefficients and standard errors for which are reported in Column (2) of Table 2. The y-axis reports the quantity of sales not explained by our econometric model and the x-axis reports the date (i.e., month of sample). Each (blue) dot is a monthly observation and the (red) reference line is aligned with the month the regulation is implemented (July 2007).

It is clear from the residuals plot that there was some anticipation of the policy. During the months immediately prior to the implementation of the regulation the residuals are strongly and consistently positive. Immediately following the implementation, residuals are strongly and consistently negative. As our model predicts, these effects appear to be approximately symmetric. Away from the policy start date, the distribution of residuals appears quasi-random.

As an indirect test of our identifying assumption, we repeat this exercise for light-duty vehicles (LDVs), an otherwise-similar group which was not affected by the policy but which would likely have been affected by other contemporaneous macroeconomic shocks. We expect residuals for an exactly similar fixed-effects regression on LDV sales to be independent and identically distributed (i.i.d.) around the implementation of the policy. Monthly residuals for this regression are plotted in Fig. 2. From the figure, we see that these residuals are in fact i.i.d. around July 2007, suggesting that the shock which affected HDV sales around this time was not common to LDVs.

6.3. Econometric analysis of anticipation

Column (3) reports the results for a regression of the residuals from our fixed-effects model, shown in Fig. 1 above, on the two binary variables, Pre_t and $Post_t$. We confirm that the coefficients are statistically different from zero, with a strongly positive coefficient on Pre_t and strongly negative coefficient on $Post_t$. As predicted, we are unable to statistically distinguish between the absolute values of these coefficients. However, the more-negative point estimate on $Post_t$ may indicate the variable is absorbing some of the direct effect of the policy on sales.

Column (4) reports results for a model that includes Pre_t and $Post_t$ in the main regression. These results are qualitatively similar – a strongly positive coefficient on Pre_t and strongly negative coefficient on $Post_t$, with statistically indistinguishable absolute values – though this specification attenuates the coefficient on oil, and yields marginally stronger, though not statisti-

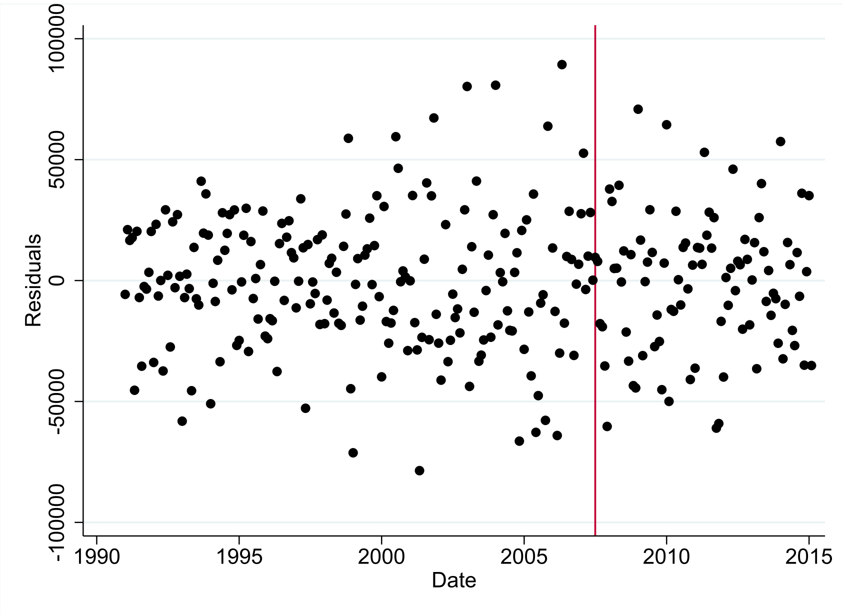

NOTE: In this figure the y-axis reports the difference between actual monthly Class-1 LDV sales and monthly Class-1 LDV sales predicted by our fixed-effects regression. The x-axis reports the date. Each dot is a monthly observation and the reference line is aligned with the month the regulation took effect (July 2007).

Fig. 2. Plot of monthly LDV sales residuals.

cally different, coefficients on both of the policy variables. That the coefficient on oil price changes suggests that oil prices covary with the treatment period. In Column (3), more of the sales variation around the time of the policy is attributed to changing oil prices, while in Column (4) Pre_t and $Post_t$ soak up more of the sales variation.

In Appendix B, we explore the sensitivity of our results to alternative specifications of the relationship between oil prices and sales, alternative specifications of the treatment period, and the inclusion of additional macroeconomic variables. We also look for evidence of similar sales shocks near the implementation of the 1998, 2002 and 2010 HDV engine standards.

7. Discussion

Our econometric analysis suggests that 31,164 more trucks were purchased ahead of regulation than would have been purchased in the absence of anticipation. This pre-buy was followed by an approximately symmetric sales slump after the policy took effect. The difference in average monthly sales during the seven months prior to and the seven months after the policy start was 11,496 vehicles. Our estimates suggest that anticipation accounts for 77% of this variation. The remainder of variation was caused by a combination of changing macroeconomic drivers, and the direct effect of the policy on sales.

We estimate how anticipation changes the net effect of the policy on emissions by breaking Equation (24) into two back-of-the-envelope calculations. The emissions increases (relative to the case of no anticipation) are determined by the difference in the lifetime emissions between the pre-bought vehicles and the counterfactual post-regulation vehicles they replace:

$$\Delta E_{(+)} = \sum_{t=1}^{T} (e_{pre_t} - e_{post_t}) VMT_t \tag{12}$$

Here, T is the vehicle lifetime in years; e_{pre_t} is the emissions rate (in grams/mile) of a pre-bought vehicle in a given year t of its life; e_{post_t} is the emission rate (in grams/mile) of a post-regulation vehicle in a given year t of its life; and VMT_t is the miles traveled per vehicle in a given year t of its life. We assume a vehicle lifetime of 10 years.[26] For emissions rates and VMT in a

[26] We take our assumed lifetime from EPA's determination of the useful life of an HDV engine (U.S. EPA, "Heavy-Duty Highway Compression-Ignition Engines and Urban Buses – Exhaust Emission Standards"). A longer lifetime would increase the emissions costs of anticipation.

Table 3
Emission impacts and economic costs due to anticipation.

Pollutant	Δ Emissions$_{(+)}$ (Thousand tons)	Δ Emissions$_{(-)}$ (Thousand tons)	Minimum Cost (Million $)	Maximum Cost (Million $)
NO_x	178.1	22.0	73.7	81.3
PM_{10}	10.2	0.6	5.0	5.2
$PM_{2.5}$	9.8	0.6	29.8	31.4
Total	–	–	108.5	117.9

given year, we use estimates from the California Air Resources Board's (CARB) Emission Factors (EMFAC) web database.[27] We compare emission rates of 2006 and 2007 model-year trucks through 2017, and (assuming that the pre-bought vehicles and counterfactual post-regulation vehicles would travel the same miles per year) use annual VMT estimates for model-year 2007 trucks.

Recall, the influx of pre-bought vehicles reduces the freight rate, pushing out the oldest vehicles with the highest operating costs, and the highest emission rates. This early retirement results in an emissions reduction (relative to the case of no anticipation):

$$\Delta E_{(-)} = Q_{ret} \times e_{ret} \times VMT \times m \tag{13}$$

Here, Q_{ret} is the quantity of vehicles which were pushed into early retirement by anticipation; e_{ret} is the emission rate of the average retired vehicle; *VMT* is the miles which the vehicle would have traveled per day for the remainder of its lifetime in the no anticipation case; and m is the average difference in the lifetime of the retired vehicles between the anticipation and no anticipation cases. As we don't know the shape of the supply curve for HDVs, we cannot say how many vehicles were pushed into early retirement by the pre-buy. From our theory model we know that the number of vehicles which leave the market early cannot exceed the number of pre-bought vehicles, so we test the two boundary cases: a minimum of zero early retirements (i.e., the case of a perfectly inelastic supply curve) and a maximum of 31,164 early retirements (i.e., the case of a perfectly elastic supply curve). The first case is simple: zero early retirements implies no emissions reduction from anticipation. For the second case, we assume that the retired vehicles were 10 years old at the time of retirement (i.e., model year 1997) and obtain emission rates and VMT for this model year in 2007 from the EMFAC web database. Since we find that the pre-buy lasted seven months on either side of the policy, we assume that the average retired vehicle left the market 7 months earlier than it would have in the case of no anticipation.

We report the results from our back-of-the envelope calculations in Table 3, below. We take social cost estimates from the Air Pollution Emission Experiments and Policy Analysis (APEEP) model, and discount the damages associated with emissions in later years at an annual rate of 5%.[28] Combined, the economic costs of anticipation are estimated to be between $108 and $118 million in 2007 dollars.

8. Conclusion

In this paper, we address four specific questions: How does anticipation of a discrete change in regulation affect the pattern of new-truck sales? How does this change in the pattern of new-truck sales affect the pattern of used-truck retirements? How do these changes in purchasing and retirement patterns affect the environmental benefits of standards? Empirically, have recent discrete changes in regulation caused firms to pre-buy trucks? To answer these questions, we first develop a theoretical model, which incorporates the effects of anticipation on new-vehicle sales and the used-vehicle fleet, and differentiates those impacts from the previously identified direct and Gruenspecht effects. We

[27] The EMFAC emissions model is developed and used by CARB to assess emissions from on-road vehicles including cars, trucks, and buses in California, and to support ARB's regulatory and air quality planning efforts to meet the Federal Highway Administration's transportation planning requirements. USEPA approves EMFAC for use in State Implementation Plan and transportation conformity analyses. The EMFAC web database reports daily VMT, running emission rates and daily idling emissions of NO_x, $PM_{2.5}$ and PM_{10} for a representative Heavy Duty Tractor Truck (vehicle category T7 tractor) of a given model year in a given calendar year in the state of California. We assume that the trucks in California are representative of the nation's trucks, and find that EMFAC's reported average VMT per vehicle per year is comparable to national average VMT for class-8 trucks. To obtain an estimate for combined running and idling emissions in g/mile, we add running emissions (g/mile) to the reported idling emissions (g/vehicle/day) divided by the reported VMT (miles/vehicle/day). Though the difference in emission rates between model years is significantly higher than the regulated change in engine emission standards, we believe EMFAC provides the measures. Emission standards are essentially a measure of the maximum efficiency achieved by one (major) vehicle component in an optimal controlled environment. In contrast, EMFAC estimates *real-world* tailpipe emission rates from samples taken during on-road testing. As we are estimating the full unintended effects of anticipation, we use the whole vehicle emissions for our calculations. For comparison, the NO_x engine standards for model-year 2006 and 2007 HDVs were 0.87 g/mile and 0.072 g/mile, respectively (using a conversion factor of 2.763 bhp-hr/mile from U.S. EPA, "Update Heavy Duty Emission Conversion Factors for MOBILE6"). The reported vehicle NO_x emissions for model years 2006 and 2007 in 2007 were 10.16 and 8.95 g/mile, respectively (including running and idling emissions as estimated by EMFAC).

[28] The reported average marginal damage is $2200/ton $PM_{2.5}$/year, $350/ton PM_{10}/year and $300/ton NO_x/year. (Muller and Mendelsohn, 2007). These are in 2002 dollars, and we convert damages to 2007 dollars using the Bureau of Labor Statistics CPI data.

test our predictions using a data set of monthly U.S. sales of new freight trucks around the time of EPA's 2007 implementation of HDV criteria pollutant standards, widely regarded as the most significant action taken by EPA (i.e., with respect to trucks) during the 25-year span of our data. Consistent with our predictions, we find evidence that anticipation caused a sales spike in the months before the policy took effect and a sales slump after implementation. We estimate anticipation of the standards caused several thousand more trucks to be sold in each of the months prior to, and approximately the same number fewer trucks to be sold in each of the months after, the introduction of the standards.

Our results have important implications for policy design and program evaluation. Ex ante, policy-makers should account for the effects of anticipation, and minimize the costs associated with it. Tradable performance standards that allow intertemporal banking and borrowing, like the current Corporate Average Fuel Economy standards for passenger cars and light-duty trucks, diminish the incentive to pre-buy by effectively phasing standards in over time.[29] Ex post, analysis that does not account for anticipation risks mischaracterizing the effects of policy, including its costs in terms of lost production and employment, and its benefits in terms of effectiveness. For analysts using time-series variation to study the effects of standards, failing to account for anticipation likely results in significantly biased estimates. In particular, confounding the effects of anticipation with the direct effects of policy would, under a variety of identification strategies, result in biased estimates of the direct effects of the policy. More broadly, our findings have important implications for the analysis of markets in which agents can shift the timing of purchases in anticipation of new regulation. Anticipation is not unique to emissions standards in the HDV industry; whenever regulation is expected to result in a discontinuous change, and agents affected by the regulation are able to adjust the timing of their behavior, we should expect to see some form of anticipation.

Although EPA has updated HDV emission standards at several other times (i.e., 1998, 2002 and 2010) covered in our data, the evidence of sales shocks around these implementations is mixed. Consistent with Lam and Bausell (2007), we find evidence of a modest pre-buy ahead of the 2002 standards – the second most-costly implementation in our time series – but not ahead of the 1998 or 2010 standards. There are several reasons firms may not have pre-bought in anticipation of these regulations. The 1998 and 2010 implementations of emission standards allowed OEMs to deploy technology that was already available, reducing the uncertainty around the reliability and operating costs of policy-compliant vehicles. In contrast, the 2002 and 2007 standards were costlier for OEMs to comply with, and the technology required to comply with these standards was, in many cases, unfamiliar to trucking firms. The uncertainty around the new technology may have strengthened the incentive to purchase pre-policy vehicles, which may have been perceived to be more reliable.

In addition to criteria pollutants, EPA regulates GHG emissions. EPA and the National Highway Transportation Safety Administration have recently finalized Phase 2 GHG and fuel efficiency standards for model-years 2019–2027, which standards are expected to significantly increase the upfront costs of new vehicles. Given the pre-buy observed in response to criteria pollutant standards, one might predict a similar response to these Phase 2 standards. However, these standards are not likely to create conditions that would induce a pre-buy. While the societal benefits of criteria pollutant standards diffuse throughout the population, the bulk of the benefits of GHG standards accrue to vehicle operators. EPA projects, within several years, the proposed Phase 2 GHG standards will fully compensate HDV operators for the increased upfront cost, and that the lifetime savings accruing to those operators will total $170 billion (U.S. EPA, "Cutting Carbon Pollution, Improving Fuel Efficiency, Saving Money and Supporting Innovation for Trucks"). In the context of our theory model, decreased lifetime operating costs, which outweigh the increase in upfront costs, would eliminate the incentive to pre-buy vehicles.

Future work may consider the implications of alternative market structures. We model competitive suppliers in a competitive freight market. Results of a model with less competition in either sector would surely yield qualitatively different results. In particular, less competition upstream would add endogeneity to the upfront price, and less competition downstream would alter the entry and exit conditions in the freight market. In addition, it would be worthwhile to verify our model empirically in a setting where macroeconomic conditions are more stable around the time of the policy change, and in a setting where data and a clear counterfactual would allow for an identification of the direct effect of the policy.

Appendix A

Equilibrium comparative statics for a change in *M*

Equilibrium M can be described in terms of equilibrium price P, and lifetime T (itself a function of P):

$$M(T,P) = \int_0^T (P(Q(t)) - C(t))e^{-rt}dt. \tag{14}$$

[29] In particular, CAFE standards phase-in over multiple years and allow firms to bank and borrow credits within multi-year compliance periods. As a result, opportunity costs of emissions reductions may be somewhat stable over time and firms may be able to arbitrage differences in compliance costs across years, reducing the total cost of compliance, and mitigating the risk of policy-induced shocks.

Taking the total derivative of $M(T, P)$ above yields:

$$dM = \frac{\partial M}{\partial T}dT + \frac{\partial M}{\partial P}dP. \tag{15}$$

The partial derivative of M with respect to T, $\frac{\partial M}{\partial T}$, is simply the change in NPV of a vehicle evaluated at time T:

$$\frac{\partial M}{\partial T} = (P(T) - C(T))e^{-rT}. \tag{16}$$

From the retirement condition, one can see that, as a consequence of the envelope theorem, $\frac{\partial M}{\partial T}$ reduces to zero. In order to take the partial with respect to P, $\frac{\partial M}{\partial P}$, we first express $M(T, P)$ as:

$$M(T, P) = \int_0^T P(t)e^{-rt}dt - \int_0^T C(t)e^{-rt}dt. \tag{17}$$

In equilibrium, the price P is constant from time $t = 0$ to T. Thus, we view $P(t)$ as a constant P, and can integrate as follows:

$$M(T, P) = \frac{-Pe^{-rt}}{r}\Big|_{t=0}^{t=T} - \int_0^T C(t)e^{-rt}dt \tag{18}$$

$$M(T, P) = \frac{-Pe^{-rT}}{r} + \frac{P}{r} - \int_0^T C(t)e^{-rt}dt \tag{19}$$

$$M(T, P) = \frac{P}{r}(1 - e^{-rT}) - \int_0^T C(t)e^{-rt}dt. \tag{20}$$

Now it is simple to take the partial derivative:

$$\frac{\partial M}{\partial P} = \frac{1 - e^{-rT}}{r}. \tag{21}$$

Which we plug into Equation (15) and rearrange to find:

$$\frac{dP}{dM} = \frac{r}{1 - e^{-rT}}. \tag{22}$$

A positive vehicle lifetime T and discount rate r require that $-rT < 0$, and thus $e^{-rT} < 1$ and $\frac{\partial P}{\partial M} > 0$.

Analysis of a change in the cost function

Without loss of generality, suppose regulation decreases the cost function of new vehicles purchased after regulation, without affecting the purchase price. In particular, let $C_2(t) \le C_1(t)$ for all $t > 0$, where $C_1(t)$ is the cost function of new vehicles purchased prior to the regulation taking effect and $C_2(t)$ the cost function of new vehicles purchased afterward. How will this change affect P, Q, and T?

A change in $C(t)$ directly affects both the entry and exit conditions for vehicles in this market. Beginning with the entry condition, a lower $C(t)$ implies a new vehicle would earn positive lifetime profit at the existing equilibrium freight rate P_1 and new-vehicle purchase price M:

$$\int_0^{T_1} (P_1 - C_1(t))e^{-rt}dt = M < \int_0^{T_1} (P_1 - C_2(t))e^{-rt}dt. \tag{23}$$

The entry condition requires that firms will purchase new vehicles until the point where the NPV of a new vehicle subject to the regulation is exactly equal to the purchase price of a new vehicle. The exit condition requires that the operating cost of a vehicle at the end of its lifetime is equal to the market freight rate. We assume no change in purchase price; thus, the available adjustment margins for meeting these conditions are vehicle lifetime, and market freight rate. Note, there is no change in lifetime T alone that, holding P steady, will satisfy both the entry and exit conditions for an equilibrium. Holding P steady at P_1, any increase in lifetime would only increase the NPV of a new vehicle. Alternatively, a decrease in T, holding P constant, would not satisfy the exit condition: $C_2(T_2) < C_1(T_2) = P_1$. Thus, it must be the case that such a decrease in the cost function results in a lower equilibrium freight rate and, fixing demand, a higher quantity of vehicles operating in the market.

The effect of the decrease in the cost function on the lifetime T of a vehicle depends on the new shape of the cost function. In the (special) linear case, where the change to $C(t)$ is simply a decrease in slope, lifetime is extended (see Fig. A1 for a representation of the change in steady state for a linear $C(t)$). To see this, recall that M must be equal to the discounted value of operating profits over a vehicle's lifetime. If the change in $C(t)$ is simply a change in the slope of a line, M can be represented by the area of the triangle below the equilibrium price and above $C(t)$. In order for M to stay constant as the slope of the cost function declines, T must increase as P falls. In other cases, T may increase or decrease in response to a change in $C(t)$, depending on the relative levels and rates of change of the old and new cost functions.

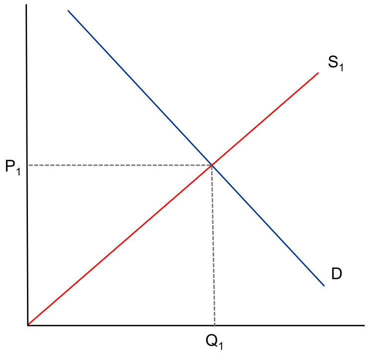

(a) Original equilibrium

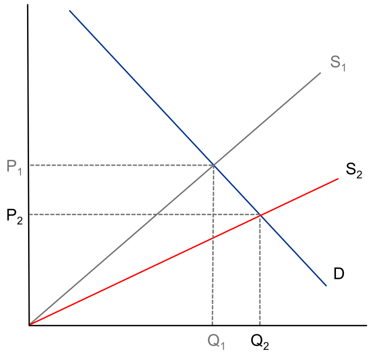

(b) New equilibrium

NOTE: In these figures the y-axis measures the freight rate, P, and the x-axis measures the quantity of vehicles in service, Q, or equivalently the quantity of freight services provided. In the first figure, the supply curve of freight services (red), which corresponds to the aggregate cost function of all trucks, slopes up from the origin to intersect the demand curve (blue) for freight services at the equilibrium P_1 and Q_1. In the second figure, the supply curve (red) has pivoted downward to intersect the demand curve (blue) at the new equilibrium P_2 and Q_2. This supply curve corresponds to the truck fleet which is in service when the first regulation-compliant truck is retired.

Fig. A1. The equilibrium effects of a cost change.

Emissions impact of anticipation

When firms change their purchasing and retirement decisions in anticipation of impending regulation, they impact the environmental outcome of that regulation. Without loss of generality, we continue to examine the case of an increase in purchase price, that is, where anticipation results in a pre-buy. Anticipation affects emissions through two channels – pre-bought vehicles create additional emissions, compared to the policy-compliant vehicles they displace, but emissions are avoided from those vehicles pushed into early retirement by the pre-buy. The net environmental impact depends on the relative size of these countervailing effects. The effect on emissions, evaluated over the lifetime of the pre-bought vehicles, is captured in the following equation:

$$\Delta E = \int_0^{T_2} (e_{pre}\overline{Q}_{pre} - e_{post}Q_{post}(t) - e_{ret}(\overline{Q}_{ret} - Q_{ret}(t)))dt. \tag{24}$$

In discussing the emissions impacts, we will use "Base Case" to refer to the scenario without anticipation, and "Anticipation Case" to refer to the scenario in which firms anticipate future regulation. We define e_{pre} as the average emission rate for a pre-bought vehicle, e_{ret} as the average emission rate for a vehicle which is retired before the market reaches the equilibrium in the Anticipation Case, and e_{post} as the average emission rate for a vehicle which is purchased after the policy is implemented. Finally, let $\overline{Q}_{pre}$ be the magnitude of the pre-buy, $\overline{Q}_{ret}$ be the magnitude of vehicles pushed into early retirement by the pre-buy, $Q_{post}(t)$ be the difference in the quantity of vehicles purchased after the policy is implemented in the Base and the Anticipation cases, and $Q_{ret}(t)$ be the difference in the quantity of vehicles retired after the policy is implemented in the Base and Anticipation cases.[30]

It is instructive to evaluate the environmental outcome in three distinct periods, during which anticipation affects emissions, and between which that effect differs. Let t_0 be the time of the pre-buy (before policy is implemented), t_1 be time when the market reaches equilibrium (Q_2, P_2) in the Base Case, t_2 be the time when the market reaches equilibrium (Q_2, P_2) in the Anticipation Case (τ in Equations (6) through (8) above), and t_3 be the time when the pre-bought vehicles exit the market in

[30] Note that $\overline{Q}_{pre}$ refers to the quantity of vehicles which were purchased prior to the implementation, as a result of expectations about the regulation. This does not include those vehicles which would have been purchased in the absence of any anticipation of policy changes. Similarly, $\overline{Q}_{ret}$ refers only to the quantity of vehicles which are retired early as an immediate result of the pre-buy. We have used $\overline{Q}$ to refer to a quantity of vehicles which remains constant over time, while $Q(t)$ refers to a quantity of vehicles which changes over time.

the Anticipation Case. To discuss the dynamics of emissions effects, we will consider, as an illustrative example, the special case where the slopes of the supply and demand curves are constant. Constant slopes imply that vehicles are retired at the same rate in the Anticipation Case, between t_0 and t_2, as they are in the Base Case between t_0 and t_1.

Between t_0 and t_1, no vehicles are purchased in either the Base Case or the Anticipation Case (i.e. $Q_{post}(t) = 0$). The retirement rate is the same in both cases (i.e. $Q_{ret}(t) = 0$). All emissions from pre-bought vehicles in the Anticipation Case are additional, while emissions from the quantity of vehicles pushed into early retirement in the Anticipation Case are avoided. During this period, the emissions impact of the anticipation is:

$$\Delta E_{0-1} = (t_1 - t_0)(e_{pre}\overline{Q}_{pre} - e_{ret}\overline{Q}_{ret}). \tag{25}$$

Between t_1 and t_2, the market is in equilibrium in the Base Case, but has not yet reached equilibrium in the Anticipation Case. As vehicles are purchased in the Base Case, $Q_{post}(t)$ increases, eroding the additional emissions from the pre-bought vehicles. At the same time, vehicles are retiring more quickly in the Base Case than in the Anticipation Case, increasing $Q_{ret}(t)$ and reducing the difference between total retirement in the Base and Anticipation Case, and thus reducing total avoided emissions. During this period, the emissions impact of anticipation is:

$$\Delta E_{1-2} = \int_{t_1}^{t_2} (e_{pre}\overline{Q}_{pre} - e_{post}Q_{post}(t) - e_{ret}(\overline{Q}_{ret} - Q_{ret}(t)))dt. \tag{26}$$

By time t_2, the market has reached equilibrium in both the Base and the Anticipation cases. The supply curve has shifted to the same point, crossing the demand curve at the new equilibrium. The operating cost $C(t_2)$ of the next retiring vehicle is equal to P_2, implying that all vehicles with $C(t_2) > P_2$ have already been retired in both cases. Thus, the total quantity of retired vehicles in the Base Case is now equal to that in the Anticipation Case (i.e. $Q_{ret}(t) = 0$), and the total quantity of purchased vehicles in each case must also be equal. Since no vehicles have been purchased post-policy in the Anticipation Case before t_2, we have $Q_{post}(t_2) = \overline{Q}_{pre}$. From t_2 to t_3, vehicles are purchased and retired at the same rate in both cases. The only difference in emissions stems from the higher emission rate of the pre-bought vehicles compared to that of the policy-compliant vehicles purchased in the Base Case. During this period, the emissions impact of anticipation is:

$$\Delta E_{2-3} = (t_3 - t_2)(e_{pre} - e_{post})\overline{Q}_{pre}. \tag{27}$$

Appendix B

Robustness Checks

Sensitivity to treatment-period specification

Table B1 presents results for our fixed-effects model, including Pre_t and $Post_t$, across a range of treatment-period specifications. Column (4) corresponds to our central specification (also Column (4) in Table 2).

Table B1
Robustness to Treatment-period Specification.

Variables	(1) Pre: 4m Post: 4m	(2) Pre: 5m Post: 5m	(3) Pre: 6m Post: 6m	(4) Pre: 7m Post: 7m	(5) Pre: 8m Post: 8m
Real Oil Price	−33.58	−19.26	−20.62	−28.19	−28.35
	(12.02)	(11.01)	(10.94)	(11.38)	(11.59)
GDP	3.925	3.632	4.224	5.018	5.408
	(1.363)	(1.231)	(1.229)	(1.295)	(1.329)
Pre-treatment	4220	2218	7474	4526	4089
	(1113)	(1185)	(1227)	(938.1)	(819.8)
Post-treatment	−6824	−9599	−3955	−5191	−4916
	(1116)	(1201)	(1070)	(862.5)	(781.7)
Observations	278	278	278	278	278
R-squared	0.898	0.917	0.918	0.910	0.907
Year FE	Y	Y	Y	Y	Y
Month FE	Y	Y	Y	Y	Y

Note: This table reports coefficients and standard errors for five specifications of our econometric model of truck sales. In each column, the dependent variable is monthly Class-8 HDV sales. Each specification includes Real Oil Price, GDP, year and month-of-year fixed effects, and two binary variables (Pre-treatment and Post-treatment). The column titles report the number of months prior to regulation for which Pre-treatment takes the value of one, and the number of months after regulation takes effect for which Post-treatment takes the value of one. Pre- and Post-treatment each take the value of zero in all other months, including the month the regulation takes effect (July 2007).

Our results are robust to alternative specifications of the pre and post period. The coefficients on Pre_t and $Post_t$ are statistically significant and generally consistent across these specifications. The five- and six-month specifications produce the most markedly different coefficients on the policy variables. In addition, these specifications break the pattern of an otherwise-stable oil-price coefficient; oil prices and policy variables drive sales to different extents in those two specifications. One explanation for the pattern of results reported in Columns (2) and (3) is that the five- and six-month specifications covary particularly closely with periods of extreme oil-price variation, entangling, in those specifications, the effect of the regulation with the effect of oil-price variation.

Sensitivity to oil-price specification

A challenge to our identification strategy is to correctly specify the relationship between the real oil price and sales. The oil price changes dramatically and unprecedentedly around the time the 2007 criteria pollutant standards were implemented (see Fig. B1). We have observed that the coefficient on oil prices is sensitive both to the addition of Pre_t and $Post_t$, and to the specification of those variables, consistent with the fact that treatment covaries with the oil prices. As a result of this covariance, it appears that in certain specifications the policy variables soak up sales variation that may, in reality, be driven by the oil price. We explore the consequences of this covariance by testing the sensitivity of our results to alternative specifications of the impact of oil price on sales, identifying the effect of the policy by regressing the sales residuals on the policy variables.

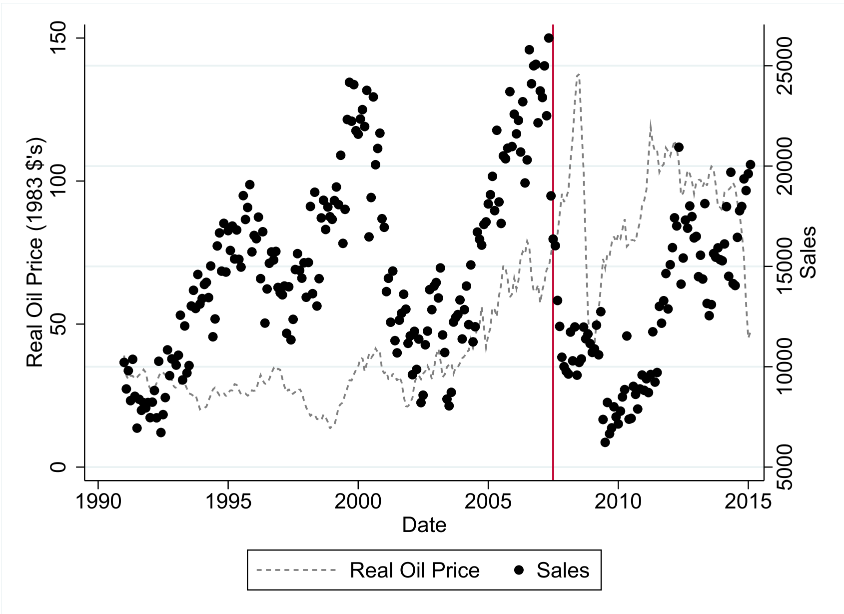

NOTE: Monthly sales are for Class-8 HDVs, and are purchased from Ward's Automotive, Inc. Oil prices are real monthly average imported crude oil prices, in 1982-1984 dollars, as reported by the Energy Information Administration (EIA). The reference line corresponds to the month the regulation took effect (July 2007)

Fig. B1. Monthly HDV sales and oil prices.

Table B2 reports coefficients and standard errors for four two-stage specifications of truck sales. Columns (1), (3), (5) and (7) each report results for a first-stage fixed-effects regression of monthly Class-8 HDV sales (excluding Pre_t and $Post_t$). Columns (2), (4), (6) and (8) report results of second-stage regressions of the residuals from the previous column's specification on the policy variables. For each of these second-stage regressions, the independent variables are two binary variables, Pre-treatment and Post-treatment, which respectively take the value of one during the seven months prior to regulation, and the seven months after regulation took effect. Columns (1) and (2) reproduce Columns (2) and (3) from Table 2.

Table B2
Sensitivity to Oil-price Specification.

Variables	(1) Sales (all)	(2) Residuals (all)	(3) Sales (Excl. 2007)	(4) Residuals (Excl. 2007)	(5) Sales (shift)	(6) Residuals (shift)	(7) Sales (shift Excl. 2007)	(8) Residuals (shift Excl. 2007)
Real Oil Price	−54.21		−9.932					
	(14.00)		(10.81)					
GDP	4.330		3.429		4.087		3.269	
	(1.610)		(1.190)		(1.599)		(1.183)	
Pre-treatment		4030		28,425		3955		28,032
		(530.0)		(1261)		(525.0)		(1245)
Post-treatment		−4676		14,804		−4624		14,565
		(530.0)		(1261)		(525.0)		(1245)
Oil (pre)					63.86		70.34	
					(52.77)		(38.95)	
Oil (post)					−60.24		−14.51	
					(14.11)		(10.93)	
Observations	278	278	266	278	278	278	266	278
R-squared	0.856	0.336	0.917	0.697	0.859	0.334	0.919	0.696
Year FE	Y	Y	Y	N	Y	N	Y	N
Month FE	Y	Y	Y	N	Y	N	Y	N

Note: This table reports coefficients and standard errors for four two-stage specifications of truck sales. Columns (1), (3), (5) and (7) each report results for a first-stage fixed-effects regression of monthly Class-8 HDV sales (excluding Pre-treatment and Post-treatment variables). Columns (2), (4), (6) and (8) report results of second-stage regressions of the residuals from the previous column's specification on Pre-treatment and Post-treatment. For each of these second-stage regressions, the independent variables are two binary variables, Pre-treatment and Post-treatment, which respectively take the value of one during the seven months prior to regulation, and the seven months after regulation took effect. Columns (1) and (2) reproduce Columns (2) and (3) from Table 2. Columns (3) and (4) report results for a specification that excludes observations from 2007 in the first-stage regression. Columns (5) and (6) report results for a specification which includes two separate independent variables for monthly real oil price before 2002, and after 2001, respectively. Columns (7) and (8) report results for a specification which both excludes observations from 2007, and includes the two separate variables for monthly real oil price in the first-stage regression.

We've seen that oil prices co-vary with the treatment period. It's possible that our estimated coefficient on oil prices in previous specifications is soaking up variation which is in fact driven by the policy variables. To test this hypothesis, we suppose that the relationship between oil prices and vehicle sales is equal in 2007 to the mean value of that relationship in all other years and calculate the implied residuals for 2007 by differencing actual and predicted monthly sales. Columns (3) and (4) report results for this specification. Somewhat surprisingly, the coefficient on oil in this specification becomes insignificant. The coefficients on Pre_t and $Post_t$ break significantly from the relatively tight range we have observed over other specifications. These results suggest that the specified effect of oil price on sales may, here, be overly restrictive. Real oil prices rose 66% in 2007, to $95/barrel, the highest observed price in our data up to that point. If firms responded to the significant rise in oil price differently than to more moderate price changes during periods of relative price stability, the relationship between oil price and truck sales during this volatile period might be importantly different.

If the effect of oil price on sales did change over time, a natural moment for that change to begin might have been during the post-2001 run up in oil prices. Between 1991 and 2001, the mean real oil price was $28/barrel, with a standard deviation of $6/barrel. Between 2002 and 2015, the mean real oil price was $74/barrel, with a standard deviation of $27/barrel. It is conceivable that this period of increasing, and increasingly volatile, oil prices marked a change in the relationship between the real oil price and HDV sales. To test this hypothesis, we interact oil price with indicator variables for pre-2002 and post-2001; Oil_{pre} takes the value of the monthly real oil price between 1991 and 2001 and zero otherwise, while Oil_{post} takes the value of the monthly real oil price between 2002 and 2015, and zero otherwise. Columns (5) and (6) report first- and second-stage regression results for this specification, respectively. Our results suggest that there was a structural shift in the relationship after 2001. The coefficient on Oil_{pre} is positive and insignificant, and the coefficient on Oil_{post} is negative and significant. The coefficients on the policy variables are statistically indistinguishable from those in Column (2). While these results suggest the effect of oil prices on sales did change significantly post-2001, as before, it may be that the coefficient on Oil_{post} is affected by sales variation that is, in fact, caused by the policy.

To address this potential source of bias, we combine the two previous specifications, estimating coefficients for Oil_{pre} and Oil_{post}, excluding data from 2007. Here, the exclusion of 2007 is likely to be somewhat less restrictive than in the previous specification. If there was a structural break, the mean response to oil price changes in 2007 is likely closer to the mean response around that time than it is to the mean response over the entire time period. Columns (7) and (8) report first- and second-stage regression results for this specification, respectively. In this specification, oil price is a positive, significant (at the 10% level) driver of sales before 2002, and a negative, insignificant driver after 2001 (though this point estimate is not statistically different from the point estimate in Column (4) of Table 2). The estimated value of Pre_t and $Post_t$ fall within the range of estimates observed in Table 2.

While we are unable to definitively disentangle the effects of anticipation from the (changing) effect of oil prices, we find that the estimated impact of anticipation is relatively stable over a range of plausible assumptions.

Sensitivity to macroeconomic variables

Some readers may be particularly concerned that the policy shortly preceded the "Great Recession". If macroeconomic sales drivers, (e.g., expectations about the future demand for freight activity) or expectations about future oil prices were changing discretely around the time of the policy, our coefficients may be biased. We attempt to address these concerns with seven additional model specifications, which include the level of the S&P 500 stock market index, the U.S. Treasury securities yield curve (i.e., the difference in effective interest rates between the 3-month and 10-year securities), the six-month London Interbank Offer Rate (LIBOR), the 3-month lead of GDP, and the 3-month, 6-month and 12-month oil futures contracts traded on the New York Mercantile Exchange (NYMEX), respectively.[31] If sharp changes in any of these leading economic indicators were driving changes in sales around the time the policy took effect, including the indicator as a covariate should attenuate our policy coefficients. Table B4 reports results for these models. For reference, Column (1) reports results for our preferred specification. The policy coefficients are quite stable and remain statistically robust across these alternative specifications, suggesting that omitting these variables from our preferred specification is not biasing our policy coefficients.

Table B3
Descriptive Statistics of Additional Macroeconomic Variables.

Variable	Obs	Mean	Std. Dev.	Min	Max
S&P 500 Index	278	1.866	1.187	−0.697	3.685
Yield Curve (3mo-10yr)	278	3.335	2.19	0.323	6.975
LIBOR (6-Mo)	278	1096.291	397.227	407.36	2082.2
GDP Lead (3-Mo)	278	11842.66	3391.777	6380.8	17703.7
Oil Futures (3-Mo)	278	51.205	33.958	12.69	125.1
Oil Futures (6-Mo)	278	51.335	32.676	13.77	117.31
Oil Futures (12-Mo)	278	50.982	30.646	14.36	107.24

Table B4
Sensitivity to macroeconomic variables.

Variables	(1) Sales	(2) Sales	(3) Sales	(4) Sales	(5) Sales	(6) Sales	(7) Sales	(8) Sales
Real Oil Price	−28.19	−22.02	−24.10	−27.62	−25.12	27.85	15.17	−32.96
	(11.38)	(11.83)	(10.62)	(10.75)	(12.60)	(39.86)	(37.81)	(31.71)
GDP	5.018	6.192	2.779	2.663	4.472	5.173	5.207	4.984
	(1.295)	(1.444)	(1.260)	(1.296)	(1.516)	(1.296)	(1.303)	(1.315)
Pre-treatment	4526	4795	3430	4278	3793	4561	4595	4508
	(938.1)	(945.6)	(891.7)	(886.9)	(986.8)	(936.2)	(939.0)	(947.0)
Post-treatment	−5191	−5007	−5422	−5129	−5079	−5074	−5087	−5197
	(862.5)	(864.4)	(804.0)	(814.3)	(865.9)	(864.1)	(866.0)	(865.0)
S&P 500 Index		−2.774						
		(1.537)						
Yield Curve (3mo-10yr)			−1224					
			(199.5)					
LIBOR (6-Mo)				1145				
				(208.6)				
GDP Lead (3-Mo)					3.487			
					(1.404)			
Oil Futures (3-Mo)						−68.31		
						(46.59)		

(continued on next page)

[31] Daily levels of the S&P 500 stock market index were collected from the Chicago Board of Exchange (CBOE). Daily LIBOR rates and treasury yields were collected from the Federal Reserve Bank of St. Louis. Daily oil futures contracts were collected from NYMEX. We use simple monthly averages of these variables in our regressions. Table B3 reports summary statistics for these variables.

Table B4 *(continued)*

Variables	(1) Sales	(2) Sales	(3) Sales	(4) Sales	(5) Sales	(6) Sales	(7) Sales	(8) Sales
Oil Futures (6-Mo)							−62.34	
							(51.83)	
Oil Futures (12-Mo)								9.371
								(58.13)
Observations	278	278	278	278	267	278	278	278
R-squared	0.910	0.911	0.922	0.920	0.912	0.911	0.910	0.910
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Month FE	Y	Y	Y	Y	Y	Y	Y	Y

Note: This table reports coefficients and standard errors for eight specifications of our econometric model of truck sales. In each column, the dependent variable is monthly Class-8 HDV sales. Each specification includes Real Oil Price, GDP, year and month-of-year fixed effects, and two binary variables, Pre-treatment and Post-treatment, which respectively take the value of one during the seven months prior to regulation, and the seven months after regulation took effect. Column (1) replicates our preferred specification. Columns (2)–(8) include each of the variables in our preferred specification, plus the level of the S&P 500, the yield curve, the six-month LIBOR, the three-month lead of GDP, and the 3-month, 6-month and 12-month NYMEX oil futures, respectively.

Evidence of other pre-buys

During the span of our data, EPA implemented four rounds of new-HDV criteria pollutant standards, for model year 1998, 2002, 2007 and 2010 vehicles, respectively.[32] While we concentrate our analysis on the 2007 update to emission standards (widely regarded as the most significant of the four), we have tested for anticipation around each of the standards. Table B5 reports results of regressions which replicate our preferred specification for each of the regulations. Column (1) replicates Column (4) in Table 2. Columns (2) through (4) include additional binary treatment variables, corresponding to the seven months ahead of and following implementation of the 1998, 2002 and 2010 standards, respectively.

We find evidence of a pre-buy ahead of the 2002 standards, followed by an approximately symmetric sales slump in the months after implementation. We do not find similar evidence of anticipation around the 1998 and 2010 standards. There are several reasons firms may not have pre-bought in anticipation of these regulations. OEMs could comply with the 1998 and 2010 emission standards by deploying technology that was already available, reducing the uncertainty around the reliability and operating costs of policy-compliant vehicles. In contrast, the 2002 and 2007 standards were both costlier for OEMs to comply with, and required them to introduce new and unfamiliar technology. Trucking firms may have had concerns around the reliability and operating costs of new technologies, which uncertainty may have strengthened the incentive to purchase pre-policy vehicles.

Table B5
Evidence of other pre-buys.

Variables	(1) Sales	(2) Sales	(3) Sales	(4) Sales
Real Oil Price	−28.19	−28.80	−38.24	−28.60
	(11.38)	(11.42)	(10.80)	(11.35)
GDP	5.018	5.041	6.617	5.043
	(1.295)	(1.299)	(1.248)	(1.304)
Pre-treatment	4526	4467	4463	4473
	(938.1)	(940.0)	(879.3)	(930.7)
Post-treatment	−5191	−5237	−5224	−5266
	(862.5)	(863.9)	(808.1)	(855.9)
Pre-treatment (1998)		−1008		
		(924.1)		
Post-treatment (1998)		−1027		
		(853.6)		
Pre-treatment (2002)			2627	
			(797.7)	

(*continued on next page*)

[32] Note, the regulation which we refer to as the 2002 model-year standard originally applied to model-year 2004 vehicles. However, EPA accelerated the implementation of these standards, such that in fact they applied to vehicles sold after October 2002. Model-year 2010 vehicles were covered under the same regulation as 2007-model year vehicles, but were subject to distinctly different requirements for compliance.

Table B5 *(continued)*

Variables	(1) Sales	(2) Sales	(3) Sales	(4) Sales
Post-treatment (2002)			−2018	
			(1024)	
Pre-treatment (2010)				−693.4
				(928.8)
Post-treatment (2010)				−1960
				(845.2)
Observations	278	278	278	278
R-squared	0.910	0.911	0.922	0.912
Year FE	Y	Y	Y	Y
Month FE	Y	Y	Y	Y

Note: This table reports coefficients and standard errors for four specifications of our econometric model of truck sales. In each column, the dependent variable is monthly Class-8 HDV sales. Each specification includes Real Oil Price, GDP, year and month-of-year fixed effects, and two binary variables, Pre-treatment and Post-treatment, which respectively take the value of one during the seven months prior to regulation, and the seven months after regulation took effect. Columns (2)–(4) each include two additional binary treatment variables corresponding to the seven months prior to and following the implementations of the model-year 1998, 2002 and 2010 HDV engine standards, respectively.

Table B6
Sensitivity to lagged variables.

Variables	(1) Sales	(2) Sales	(3) Sales	(4) Sales	(5) Sales	(6) Sales	(7) Sales	(8) Sales	(9) Sales
Real Oil Price	−28.19	−39.00	−34.03	−29.28	−29.44	−29.16	−25.21	−19.04	−17.50
	(11.08)	(21.94)	(12.81)	(10.54)	(11.89)	(13.93)	(14.33)	(13.32)	(13.26)
GDP	5.018	5.174	5.224	5.217	5.212	5.237	5.010	2.932	2.991
	(1.215)	(1.247)	(1.247)	(1.214)	(1.221)	(1.258)	(1.299)	(1.894)	(1.653)
Pre-treatment	4526	4560	4582	4513	4419	4458	4559	4697	4714
	(2180)	(2164)	(2168)	(2201)	(2192)	(2227)	(2248)	(2298)	(2419)
Post-treatment	−5191	−5169	−5162	−5203	−5272	−5231	−5115	−5345	−5345
	(1627)	(1623)	(1624)	(1629)	(1617)	(1651)	(1689)	(1671)	(1756)
Lagged Oil (1M)		13.90							
		(25.55)							
Lagged Oil (2M)			12.33						
			(17.66)						
Lagged Oil (3M)				3.396					
				(15.44)					
Lagged Oil (4M)					−3.952				
					(15.41)				
Lagged Oil (5M)						−0.179			
						(14.29)			
Lagged Oil (6M)							8.089		
							(11.83)		
Lagged GDP (1Q)								3.204	
								(1.741)	
Lagged GDP (2Q)									3.367
									(1.239)
Observations	278	277	276	275	274	273	272	275	272
R-squared	0.910	0.910	0.909	0.909	0.908	0.908	0.907	0.910	0.909
Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Month FE	Y	Y	Y	Y	Y	Y	Y	Y	Y

Note: This table reports coefficients and standard errors for nine specifications of our econometric model of truck sales. In each column, the dependent variable is monthly Class-8 HDV sales. Each specification includes Real Oil Price, GDP, year and month-of-year fixed effects, and two binary variables, Pre-treatment and Post-treatment, which respectively take the value of one during the seven months prior to regulation, and the seven months after regulation took effect. Column (1) replicates our preferred specification. Columns (2)–(9) include each of the variables in our preferred specification, plus the one-month through six-month lag of real oil price, and the one-quarter and two-quarter lag of GDP, respectively.

Table B7
Sensitivity to leading variables.

Variables	(1) Sales	(2) Sales	(3) Sales	(4) Sales	(5) Sales	(6) Sales	(7) Sales	(8) Sales	(9) Sales
Real Oil Price	−28.19	−23.60	−29.67	−31.25	−29.90	−27.64	−26.84	−36.70	−37.13
	(11.08)	(27.93)	(16.93)	(14.28)	(13.09)	(12.65)	(12.53)	(14.27)	(14.95)
GDP	5.018	5.049	4.997	4.802	4.706	4.778	4.966	0.825	2.298
	(1.215)	(1.203)	(1.205)	(1.239)	(1.274)	(1.302)	(1.353)	(1.584)	(1.487)
Pre-treatment	4526	4535	4518	4512	4458	4416	4485	3518	2820
	(2180)	(2212)	(2183)	(2156)	(2144)	(2148)	(2173)	(2003)	(2026)
Post-treatment	−5191	−5159	−5220	−5387	−5595	−5701	−5478	−5900	−6603
	(1627)	(1684)	(1656)	(1632)	(1615)	(1659)	(1679)	(1400)	(1492)
Lead Oil (1M)		−4.766							
		(23.69)							
Lead Oil (2M)			1.855						
			(12.99)						
Lead Oil (3M)				8.074					
				(10.66)					
Lead Oil (4M)					11.45				
					(10.46)				
Lead Oil (5M)						11.16			
						(11.73)			
Lead Oil (6M)							5.938		
							(13.07)		
Lead GDP (1Q)								5.474	
								(1.875)	
Lead GDP (2Q)									5.327
									(2.132)
Observations	278	277	276	275	274	273	272	275	272
R-squared	0.910	0.910	0.909	0.909	0.909	0.909	0.908	0.914	0.914
Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Month FE	Y	Y	Y	Y	Y	Y	Y	Y	Y

Note: This table reports coefficients and standard errors for nine specifications of our econometric model of truck sales. In each column, the dependent variable is monthly Class-8 HDV sales. Each specification includes Real Oil Price, GDP, year and month-of-year fixed effects, and two binary variables, Pre-treatment and Post-treatment, which respectively take the value of one during the seven months prior to regulation, and the seven months after regulation took effect. Column (1) replicates our preferred specification. Columns (2)–(9) include each of the variables in our preferred specification, plus the one-month through six-month lead of real oil price, and the one-quarter and two-quarter lead of GDP, respectively.

Appendix C

Supplementary data

Supplementary data related to this article can be found at https://doi.org/10.1016/j.jeem.2018.03.005.

References

40 C.F.R. § 69, 80, 86. Control of air pollution from new motor vehicles: heavy-duty engine and vehicle standards and Highway diesel fuel sulfur control requirements.
Calpin, Patrick, Plaza-Jennings, Esteban, 2012. A Look Back at EPA's Cost and Other Impact Projections for MY 2004-2010 Heavy-duty Truck Emissions Standards. NADA Comments to NHTSA/EPA Re: MY 2017-2025 Proposed Standards.
Cashin, D., Unayama, T., 2016. Measuring intertemporal substitution in consumption: evidence from a VAT increase in Japan. Rev. Econ. Stat. 98 (2), 285–297.
Coglianese, John, Davis, Lucas W., Kilian, Lutz, Stock, James H., 2016. Anticipation, tax avoidance, and the price elasticity of gasoline demand. J. Appl. Econ. 32 (1), 1–15.
Crossley, T., Low, H., Sleeman, C., 2014. Using a Temporary Indirect Tax Cut as a Fiscal Stimulus: Evidence from the UK. University of Essex. Unpublished Manuscript.
D'Acunto, Francesco, Hoang, Daniel, Weber, Michael, 2016. The Effect of Unconventional Fiscal Policy on Consumption Expenditure. NBER Working Paper 22563.
Emission Rule Change the Engine for Truck Sales: But This Year's Boom to Be Next Year's bust." Chicago Tribune. Web. 23 May 2016.
Fullerton, Don, West, Sarah E., 2002. Can taxes on cars and on gasoline mimic an unavailable tax on emissions? J. Environ. Econ. Manag. 43 (1).
Gruenspecht, Howard K., 1982. Differentiated regulation: the case of auto emissions standards. Am. Econ. Rev. 72 (2), 328–331.
Harrison, David, LeBel, Mark, 2008. Customer Behavior in Response to the 2007 Heavy-duty Engine Emission Standards: Implications for the 2010 NO_x Standard. NERA Economic Consulting.
Hausman, Catherine, Rapson, David S., 2017. Regression discontinuity in time: Considerations for environmental Economists. Working Paper.
Hoekstra, Mark, Puller, Steven L., West, Jeremy, 2017. Cash for corollas: when stimulus reduces spending. Am. Econ. J. Appl. Econ. 9 (3), 1–35.

Jacobsen, Mark R., Van Benthem, Arthur A., 2015. Vehicle scrappage and gasoline policy. Am. Econ. Rev. 105 (3), 1312–1338.
Lam, Terence, Bausell, Charles, 2007. Strategic behaviors toward environmental regulation: the case of the trucking industry. Contemp. Econ. Pol. 25 (1), 3–13.
Li, Shanjun, Linn, Joshua, Spiller, Elisheba, 2013. Evaluating cash-for-clunkers: program effects on auto sales and the environment. J. Environ. Econ. Manag. 65 (2), 175–193.
Maloney, Michael T., Brady, Gordon L., 1988. Capital turnover and marketable pollution rights. J. Law Econ. 31 (1), 203–226.
McCrary, Justin, 2008. Manipulation of the running variable in the regression discontinuity design: a density test. J. Econ. 142 (2), 698–714.
Mian, Atif, Sufi, Amir, 2012. The effects of fiscal stimulus: evidence from the 2009 cash for clunkers program. Q. J. Econ. 127 (3), 1107–1142.
Muller, Nicholas Z., Mendelsohn, Robert, 2007. Measuring the damages of air pollution in the United States. J. Environ. Econ. Manag. 54 (1), 1–14.
Sallee, James, Slemrod, Joel, 2012. Car notches: strategic automaker responses to fuel economy policy. J. Publ. Econ. 96 (11–12), 981–999.
Stavins, Robert N., 2006. Vintage-differentiated environmental regulation. Stanford Environ. Law J. 25, 29.
To pre-buy or not. Fleet Owner. Web. 23 May 2016.
Torrey, W. Ford, Murray, Dan, 2014. An Analysis of the Operational Costs of Trucking: a 2014 Update. American Transportation Research Institute.
U.S. Department of Transportation Research and Innovative Technology Administration Bureau of Transportation Statistics, 2010. State Summaries: 2007 Commodity Flow Survey.
U.S. Environmental Protection Agency, 2000. Regulatory Impact Analysis: Heavy-duty Engine and Vehicle Standards and Highway Diesel Fuel Sulfur Control Requirements.